When AI turns institutional conditions into personal judgments

Human Systems Mediation and the evidence behind automated decisions

By Dr. Karine Coen-Sanchez

Research commentary | Coen-Sanchez Research & Policy Institute

Introducing Human Systems Mediation: An Invitation to Scholarly Dialogue

AI-assisted decisions depend on evidence produced within institutions. Human Systems Mediation (HSM), a framework I developed to examine the relationships among institutional governance, lived experience, interpretation, and accountability, offers a way to investigate the conditions behind that evidence. In this commentary, I apply HSM to what might be called the institutional pre-history of data: how organizational standards, encounters, and human responses become recorded and subsequently used in automated assessment. This is an invitation to scholarly dialogue and empirical investigation, not a claim that existing research has overlooked these relationships.

An employee stops contributing in meetings after repeatedly being interrupted or having their ideas dismissed. Their performance review describes them as quiet and lacking initiative. Later, an AI system uses those reviews to assess leadership potential and recommends against promotion.

The recommendation appears to confirm an individual limitation. Yet the evidence may partly record how the employee learned to navigate that workplace. Over time, the employee may also begin to say, “I am simply a quiet person.” An adaptation to institutional conditions can become both an organizational judgment and a personal explanation of identity.

This hypothetical example illustrates the question explored here: how can institutions produce the evidence through which AI subsequently classifies people, while losing sight of their own contribution to that evidence?

Human Systems Mediation (HSM) is a framework I developed to examine relationships between institutional governance and human experience. I apply it here to the institutional production of evidence used in AI-assisted decisions. This application asks how standards become authoritative, how people experience and respond to them, and whether those affected can change the terms on which they are assessed.

How institutional decisions become lived realities

In HSM, mediation refers to the social processes through which institutional arrangements become lived realities. It does not mean only resolving a dispute between parties. Formal policies operate through interpretation, professional norms, relationships, and unequal capacities to influence decisions. Institutions also allocate resources and impose material consequences; those effects remain part of the analysis.

What distinguishes Human Systems Mediation?

HSM examines not only how institutional arrangements shape human experience, but also how people respond to those arrangements and how their responses may later be interpreted, recorded, and treated as evidence about them. Its central analytical question is whether institutions recognize their own role in shaping the conduct and records they subsequently use to evaluate individuals. HSM therefore follows a connected sequence: institutional standards and authority; lived encounters and available resources; human interpretation, adaptation or resistance; the production of institutional records; and opportunities to challenge the resulting judgments and standards. These relationships draw on existing sociological traditions; the proposed contribution is to investigate them together, with particular attention to who can contest the interpretation and change the conditions. This commentary extends that inquiry to AI-assisted decisions, not to suggest that HSM is exclusively an AI theory.

My research on polite racism and cultural capital examines how racial exclusion can coexist with institutional civility. My subsequent work on Afro-Caribbean identity and resistance continues this concern with belonging and agency. HSM provides the broader framework through which I examine governance across institutional settings. Its application to AI requires investigation appropriate to that setting.

In education, the framework connects policy and governance, human experience, and technology. A university might adopt a common AI policy while students and employees encounter different expectations across departments. Investigating those differences requires attention to local interpretation, resources, and the authority to question decisions. Technical capability alone cannot explain implementation.

Neutrality and the maintenance of order

Standards do not arrive without a history. Someone defines appropriate participation, credible evidence, and professional conduct. Calling those choices neutral can make their assumptions less open to examination.

In my Policy Options article on institutional neutrality, I question how apparently impartial procedures can preserve unequal power. My concern extends to who establishes the standards and whose perspectives they reflect. In my Vivre ensemble interview, particularly at 7:46–8:29 and 25:30–26:12, I explain that becoming a formal procedure does not remove a standard's human origins or situated assumptions.

Maintaining order introduces a further question. Institutions need procedures that protect people and enable collective activity. They may also use expectations of civility and professionalism to restrict how established arrangements can be challenged. An organization can invite disagreement while permitting only forms of disagreement that leave its governing assumptions intact.

HSM therefore asks what order a standard maintains, who benefits from it, and whether affected people can revise it. A consultation offers limited influence if participants can describe their experiences but cannot challenge the institution's interpretation of them. Accountability requires reasons that can be examined and decisions that can be reconsidered, rather than reliance on an asserted absence of perspective.

When adaptation becomes a description of identity

Return to the employee who speaks less. They may initially recognize silence as protection. Repeated encounters can make that response feel habitual, and a description of conduct may become a description of the self. The institution may then interpret the same conduct as a stable personal trait.

This is a proposed process to investigate. Quietness is a valid disposition, and silence has many meanings. A person may remain silent strategically without internalizing an institutional judgment. Others resist, seek support, or become more outspoken. Researchers should ask how people understand their own responses and examine changes across time and settings.

The distinction matters for accountability. If an organization treats every recorded behaviour as an independent expression of individual ability, it may overlook how opportunities, sanctions, and recognition shaped that behaviour. Encouraging confidence alone would leave those conditions unexamined.

I use the metaphor of an internal algorithm for learned expectations and responses. It is distinct from a computational algorithm. People can reconsider these expectations and reconstruct their possibilities through reflection, relationships, and collective action. Institutional change must make that agency practicable.

What AI can reinforce

In the hypothetical promotion process, the system receives performance reviews as evidence. If those reviews omit interruptions, unequal recognition, or risks attached to disagreement, their descriptions may obscure how the workplace shaped participation. A system might process the records accurately while supporting an incomplete judgment about the employee.

The inquiry must establish whether the tool actually uses participation-related information, how that information affects its recommendation, and what reviewers do with the output. A language model summarizing reviews and a predictive model ranking candidates have different functions and require different scrutiny. Neither should be treated as possessing a human intention to exclude.

I use algorithmic polite racism here as a proposed extension of my earlier research on polite racism. The concern is not whether a machine speaks courteously, but whether apparently neutral, professional procedures give additional authority to racialized judgments formed under unequal conditions. Criteria such as initiative, civility, or executive presence may be interpreted without sufficient attention to unequal opportunities to demonstrate them or unequal consequences for disagreement. This is a proposition requiring investigation: neither an adverse automated decision nor a polite explanation alone establishes racism. Researchers would need evidence of racialized processes or effects and should consider alternative explanations.

Researchers could investigate whether racialized employees face different consequences for disagreement, whether those experiences enter performance records as deficiencies, and whether automated assessments reinforce the resulting judgments. They must also examine competing explanations and variation within groups.

The relationship to earlier scholarship

HSM's application builds on substantial prior work. Noble and Benjamin examine racial inequality and technological neutrality. Selbst and colleagues show why fairness analysis must include social context. Ahmed's work on institutional diversity and complaints examines how commitments to inclusion coexist with barriers to challenge.

The claim that institutions influence interpretation and action also has established foundations in Lipsky's account of policy implementation and organizational sensemaking. Realist evaluation examines how context and mechanisms shape outcomes, while participatory action research connects inquiry with participation and change. HSM's proposed workflow uses methods shared with these traditions.

Self-fulfilling prophecy concerns how expectations help produce their apparent confirmation. Bauer and Gill demonstrate in an experimental setting that disclosing inaccurate algorithmic scores can steer behaviour toward those assessments. Lee and colleagues show that automated and human moderation can suppress disclosures of racism, and that witnessing suppression can affect Black Americans' perceptions of belonging. These are direct precedents for connecting automated judgment, responses, and lived experience.

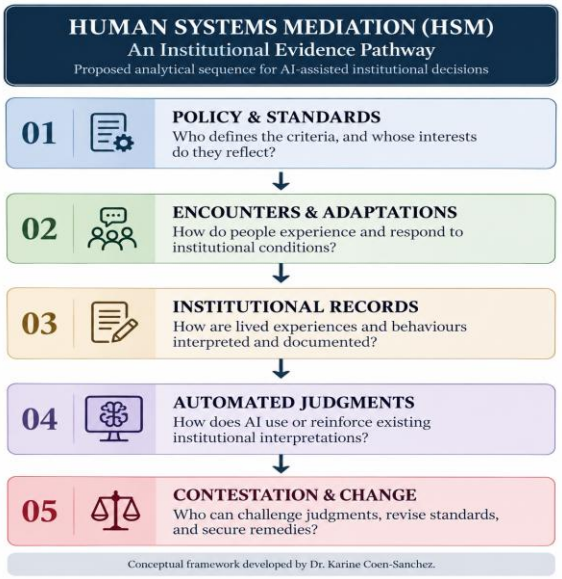
The employee example here need not begin with a prediction that the employee is quiet. It may begin with conditions that discourage speech. However, identity formation and internalization also have substantial literatures. Identity theory examines how perceived judgments enter self-understanding, while Van Dyne and colleagues distinguish motives for employee silence and the risk of observers misinterpreting them. The distinction from the classic prophecy sequence does not establish an entirely new phenomenon.

The contribution I advance is a focused application of HSM: tracing how institutional standards and practices may help produce the adaptations, records, and self-descriptions subsequently treated as evidence of personal suitability. This differs in emphasis from audits limited to statistical disparities, model outputs, or technical explanations, while remaining compatible with sociotechnical approaches that already investigate the origins and uses of data. HSM also asks who can contest the records, the standards that produced them, and the authority behind their interpretation. Its explanatory value must be assessed through comparison and evidence, rather than claims that previous scholarship ignored these relationships.

Where HSM adds to an AI assessment

An audit may ask whether outcomes differ across groups or whether a model can explain its recommendation. HSM adds questions about how the institution defined suitability, what opportunities people had before their conduct was recorded, how their responses were interpreted, and whether they can challenge both the record and the standard. These questions complement, rather than replace, established sociotechnical auditing.

Proposed analytical pathway



This sequence is a guide for inquiry, not an inevitable causal chain. People respond differently to institutional conditions; some records may be accurate and relevant, others incomplete or misleading, and AI tools can introduce additional errors of their own.

Demonstrating HSM through a practical inquiry

A focused pilot could examine one promotion process in a willing institution. I propose the following five-stage HSM inquiry as a research workflow, not a completed study, validated audit instrument, or existing software product. At each stage, the question is how institutional authority interacts with experience and with the production or contestation of evidence.

Stage 1 — Policy and interpretation mapping. Document the stated objectives, promotion criteria, decision responsibilities, and actual role of AI. Ask who defined initiative or leadership, which conduct counts as evidence, what the criteria omit, and whether affected people can question those definitions.

Stage 2 — Voluntary accounts and lived experience. Gather confidential accounts alongside relevant records, with consent and safeguards. Ask what happened before participation changed, how people understood their responses, and whether patterns differed across settings. Include contradictory accounts and, where feasible, people who disengaged. Silence should not be presumed to mean either incapacity or resistance.

Stage 3 — Evidence and data-lineage comparison. Compare accounts with reviews, source documents, and decision records. Trace when a situated experience became a written assessment, what context was included or omitted, and whether that information actually affected an AI recommendation or a human decision. Consider differences in responsibilities, resources, and preferences; no single disparity or account proves the sequence.

Stage 4 — Protected dialogue and contestation. Support opportunities for affected people to question both the accuracy of records and the criteria used to interpret their conduct. Distinguish consultation from

meaningful authority to revise standards or remedies. People should not have to disclose experiences publicly to be heard, and researchers must examine their own assumptions and positionality.

Stage 5 — Action and assessment over time. Document changes and revisit their consequences. Did the institution revise an unsupported criterion, correct a record, reconsider a decision, or reduce the risks of disagreement? Did opportunities improve? Changes in behaviour after conditions change may support a contextual interpretation, although competing explanations remain possible.

The account would need revision if the records did not influence the AI assessment, the proposed sequence did not fit participants' experiences, or another explanation better accounted for the outcome. HSM should help distinguish those possibilities, rather than interpret every finding as confirmation.

Technology supporting accountable judgment

A digital tool could organize policies, consented accounts, analytical comparisons, and agreed actions. AI could assist with theme suggestions or summaries, provided people check them against the evidence. Neither software nor an AI-generated summary should decide which experience is credible or resolve contested interpretations automatically.

The value of this application lies in following evidence back to the conditions that produced it. Institutions using AI remain responsible for the standards they authorize, the records they supply, and the decisions they make. An assessment deserves scrutiny when it turns an institutionally shaped response into a seemingly self-explanatory fact about a person.

About the author

Dr. Karine Coen-Sanchez is the founder of the Coen-Sanchez Research & Policy Institute and the developer of Human Systems Mediation. Her work examines institutional governance, racialized experiences, identity, and public policy.

Publication and AI use disclosure

This research commentary applies HSM to AI-assisted institutional assessment. The promotion example is hypothetical; the workflow is proposed for development and evaluation. This commentary is not a peer-reviewed journal article. ChatGPT assisted with research discovery, source checking, restructuring, and drafting and editing. HSM and the author's underlying research are Dr. Coen-Sanchez's contributions. Responsibility for the final published text rests with the author. Interview ideas are paraphrased from the supplied transcript; no verbatim interview quotations are used.

Copyright and scholarly use

© 2026 Dr. Karine Coen-Sanchez. All rights reserved.

Human Systems Mediation (HSM) is a theoretical framework developed by Dr. Karine Coen-Sanchez. Published by the Coen-Sanchez Research & Policy Institute (CSRPI).

Readers are welcome to cite, discuss, and critically engage with this publication with appropriate scholarly attribution. Reproduction or republication of substantial portions requires permission, subject to applicable copyright exceptions.

Coen-Sanchez, K. (2026). When AI turns institutional conditions into personal judgments: Human Systems Mediation and the evidence behind automated decisions. Coen-Sanchez Research & Policy Institute. https://www.csrpi.com/ **/hsm-commentary** Publication date: Oct 26 2026

Sources and further reading

Ahmed, S. (2012). On Being Included. Duke University Press. Publisher. Ahmed, S. (2021). Complaint! Duke University Press. Publisher.

Bauer, K., and Gill, A. (2024; first published online 2023). Mirror, Mirror on the Wall: Algorithmic Assessments, Transparency, and Self-Fulfilling Prophecies. Information Systems Research, 35(1), 226–248. DOI.

Benjamin, R. (2019). Race After Technology. Polity. Author's book page.

Coen-Sanchez, K. (2025). Polite Racism and Cultural Capital: Afro-Caribbean Negotiations of Blackness in Canada. Social Sciences, 14(8), 451. DOI.

Coen-Sanchez, K. (2026). Not Quite at Home: Afro-Caribbean Identity, Resistance, and Cultural Capital Across Generations. Genealogy, 10(1), 6. DOI.

Coen-Sanchez, K. (2026, April 15). When neutrality becomes policy, equity suffers in Canadian institutions. Policy Options. Article.

Lee, C., et al. (2024). People who share encounters with racism are silenced online by humans and machines, but a guideline-reframing intervention holds promise. PNAS, 121(38), e2322764121. DOI.

Lipsky, M. (1980; expanded edition 2010). Street-Level Bureaucracy. Russell Sage Foundation. Publisher.

Noble, S. U. (2018). Algorithms of Oppression. NYU Press. Publisher.

Pawson, R., and Tilley, N. (1997). Realistic Evaluation. Sage. Publisher. Participatory action research background: Participatory Methods.

Selbst, A. D., boyd, d., Friedler, S. A., Venkatasubramanian, S., and Vertesi, J. (2019). Fairness and Abstraction in Sociotechnical Systems. FAT*, 59–68. DOI.

Van Dyne, L., Ang, S., and Botero, I. C. (2003). Conceptualizing Employee Silence and Employee Voice as Multidimensional Constructs. Journal of Management Studies, 40(6), 1359–1392. DOI.

Weick, K. E., Sutcliffe, K. M., and Obstfeld, D. (2005). Organizing and the Process of Sensemaking. Organization Science, 16(4), 409–421. DOI.

Author interview: Vivre ensemble, supplied transcript, particularly 7:46–8:29, 12:26–13:06, 15:05–16:27, and 25:30–26:12. Video. Publication date not verified.